



The Telluride Project: History, Overview, Present Status, And A Look Ahead*

Douglas B. Kothe

Methods for Advanced Scientific Simulations Group CCS-2

**Supported by the US DOE Advanced Simulation and Computation Program*

Computer and Computational Sciences (CCS) Division
Los Alamos National Laboratory
Los Alamos, NM USA 87545
505 - 667 - 9089 (voice); 505 - 665 - 4972 (fax)
dbk@lanl.gov



Outline

Welcome

- Why this Workshop? Why are you (we) here? What are your expectations and ours?

The Telluride Project

- Motivation, charter, history, overview, current status, on the horizon
- Representative project efforts on physical models, numerical algorithms, software engineering, verification & validation, and applications
- Project team member introductions

Send Off

- Suggestions and ground rules for a mutually beneficial workshop

At Times We Work Well Together



“It’s amazing how much easier it is for a team to work together when no one has any idea where they’re going.”

And At Times We Doubt Ourselves



www.despair.com

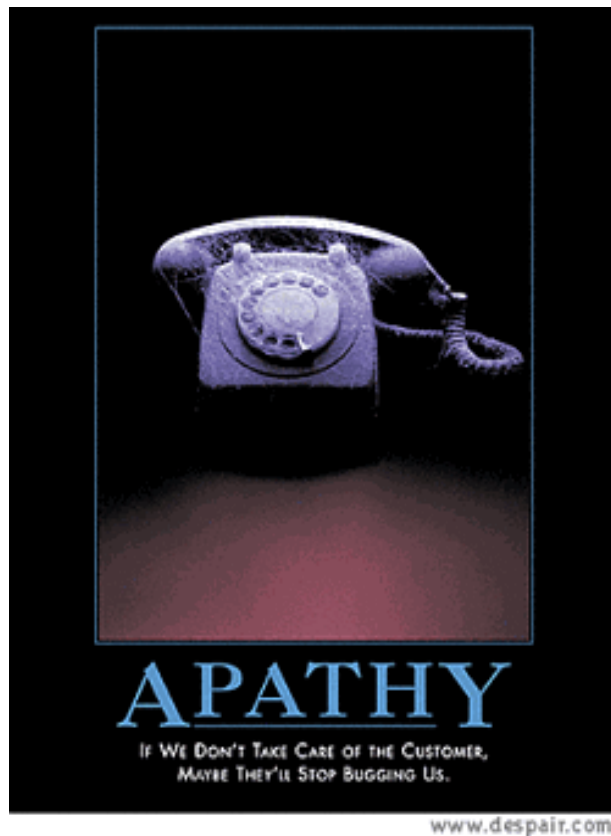
So we have meetings!

Or we

remember the mantra:

“When you start out
pathetic, there is nowhere
to go but up”

But We Always Provide Support



Workshop Overview

Motivation For The Telluride Workshop

Many government programs are justly criticized for not undergoing external scrutiny

- This scrutiny (peer review) should be self-imposed if it is not mandated by stakeholders

The DOE ASCI Program has undergone enormous scrutiny since its inception

- But mostly at higher (Lab-wide or Program Element) levels
- Project-level scrutiny has occurred sparingly and infrequently; usually only when mandated

The Telluride Project must have self-imposed peer scrutiny for a better product

We have therefore organized this workshop for three principal reasons:

1. We *cannot* declare success until or unless peer review concludes our product is value-added
2. This project is tackling a *hard* problem: we need continuous expert external advise and assistance
At a minimum, mid-course corrections; At a maximum, major project changes
3. We want to build formal collaborations with external (non-LANL) researchers such as you

We will not accomplish our workshop objectives unless we expose our weaknesses

Workshop Expectations

First, we have never done this before, so please bear with us (you are the guinea pigs)

The Telluride Project team will endeavor to

- Share all aspects of our work
 - *Physical models, numerical algorithms, software, V&V, simulation applications/analysis*
 - *The good, bad, and the ugly: false starts, mistakes, weaknesses, problems*
 - *Comparisons with other software where available and appropriate*
- Provide access to our software: So you can see and test our product for yourself

And in return we ask you to provide honest feedback (preferably written!)

- About what you like and what you do *not* like (Is this possible? Is there are better way?)
- About how and where we can improve
- On if (and how) you wish to collaborate with us

We also ask that you share your research interests and experiences with us!

How Did You Get Invited?

First, you

- are a peer whose opinion we trust and respect
- have research knowledge and experience of value to this project
- are a recognized expert in your field

Second, you might

- already be collaborating with the Telluride Project
- already possess a copy of our software (*Truchas*)
- have expressed an interest in wanting to work with us

Bottom line: we are confident your presence here will help us

The Telluride Project

Overview of the DOE ASCI Program

- which supports Telluride

Motivation, rationale, history

Physical models

Numerical algorithms

Software overview

Challenges and barriers



The DOE Advanced Simulation and Computing (ASCI) Program

Moratorium on US nuclear testing and production

- Signed into law on 1992 (Bush), extended in 1993 (Clinton)
- New weapons production has also been halted

Without underground testing (UGT), we need computer simulations to make sure our nuclear weapons stockpile is safe, reliable, and operational

- To implement these policies, the Stockpile Stewardship Program (SSP) was born

Goal: provide scientists and engineers with technical capabilities to maintain a credible nuclear deterrent without use of UGT and modernization through development of new weapons systems

- The ASCI Program was born in FY96 to achieve this goal, with a targeted completion date in FY09

Current ASCI budget: ~\$700M across DOE and several National Labs

- ~\$56M into LANL for “Advanced Application Development”, \$3.4M of which is Telluride

But how does the Telluride Project fit in?



The DOE ASCI Program

Objective: Meet the needs and requirements of the Stockpile Stewardship Program in

- *Performance:* Create predictive simulations of nuclear weapons systems to analyze behavior and assess performance in an environment without nuclear testing
- *Safety:* Predict with high certainty the behavior of full weapons systems in complex accident scenarios
- *Reliability:* Achieve sufficient, validated predictive simulations to extend the lifetime of the stockpile, predict failure mechanisms, and reduce routine maintenance → Telluride fits in here
- *Sustainability:* Use virtual prototyping and modeling to understand how new production processes and materials affect performance, safety, reliability, and aging → Telluride fits in here

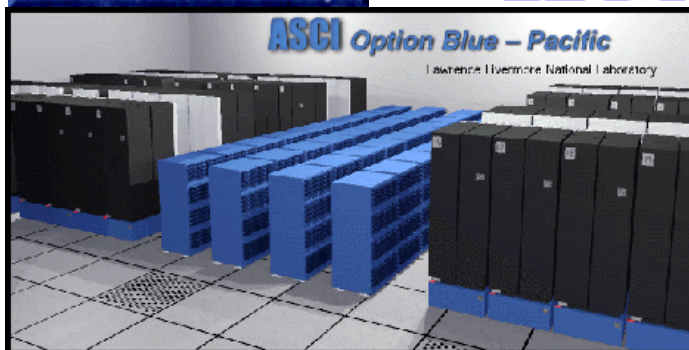
Useful ASCI URLs

- National Nuclear Security Administration (NNSA): www.nnsa.doe.gov/asc/
- Los Alamos National Laboratory (LANL): www.lanl.gov/projects/asci
- Lawrence Livermore National Laboratory (LLNL): www.llnl.gov/asci
- Sandia National Laboratory (SNL): www.sandia.gov/ASCI/index.htm

The DOE ASCI Program



SiliconGraphics



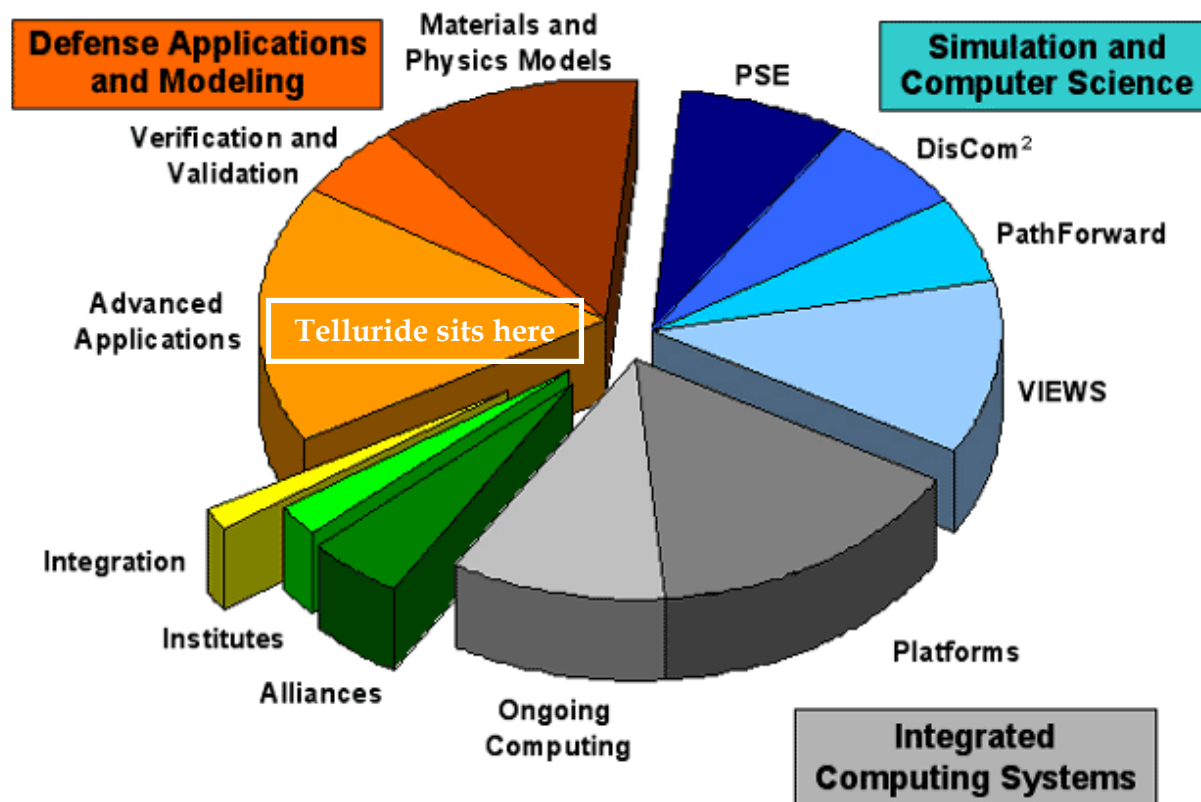
"One Program, three
Laboratories"



"Shift promptly from nuclear test-
based methods to compute-based
methods"



ASCI Program Structure



The Telluride Project

Overview of the DOE ASCI Program

- which supports Telluride

Motivation, rationale, history

Physical models

Numerical algorithms

Software overview

Challenges and barriers

Project Mission and Requirements

Deliver verified and validated computational manufacturing tools for key US Department of Energy (DOE) Complex operations

Capture, within these computational manufacturing tools:

- Realistic macroscopic models for free surface incompressible fluid flow, phase change heat transfer, material microstructure evolution, thermo-mechanical solid response, and electromagnetic effects
- Robust, high-fidelity numerical algorithms for discrete finite volume solutions of nonlinear systems of PDEs on 3D, unstructured-mesh computational domains
- Modular, extensible, maintainable, and portable software constructed with quality-assured software engineering practices

Realize efficient execution on computational platforms ranging from desktop systems to high performance computing systems typical of the DOE ASCI Program

Target casting and welding operations within the DOE Complex

Why Computational Manufacturing?

Many mission-essential manufacturing processes exist within DOE Complex

- At National Laboratories and sites such as LANL, SNL, LLNL, Y-12, KCP

New processes are needed to support future product requirements

- Re-manufacturing of Rocky Flats replacement components

These processes are unique

- Involving specialized metal alloys and glove box operations

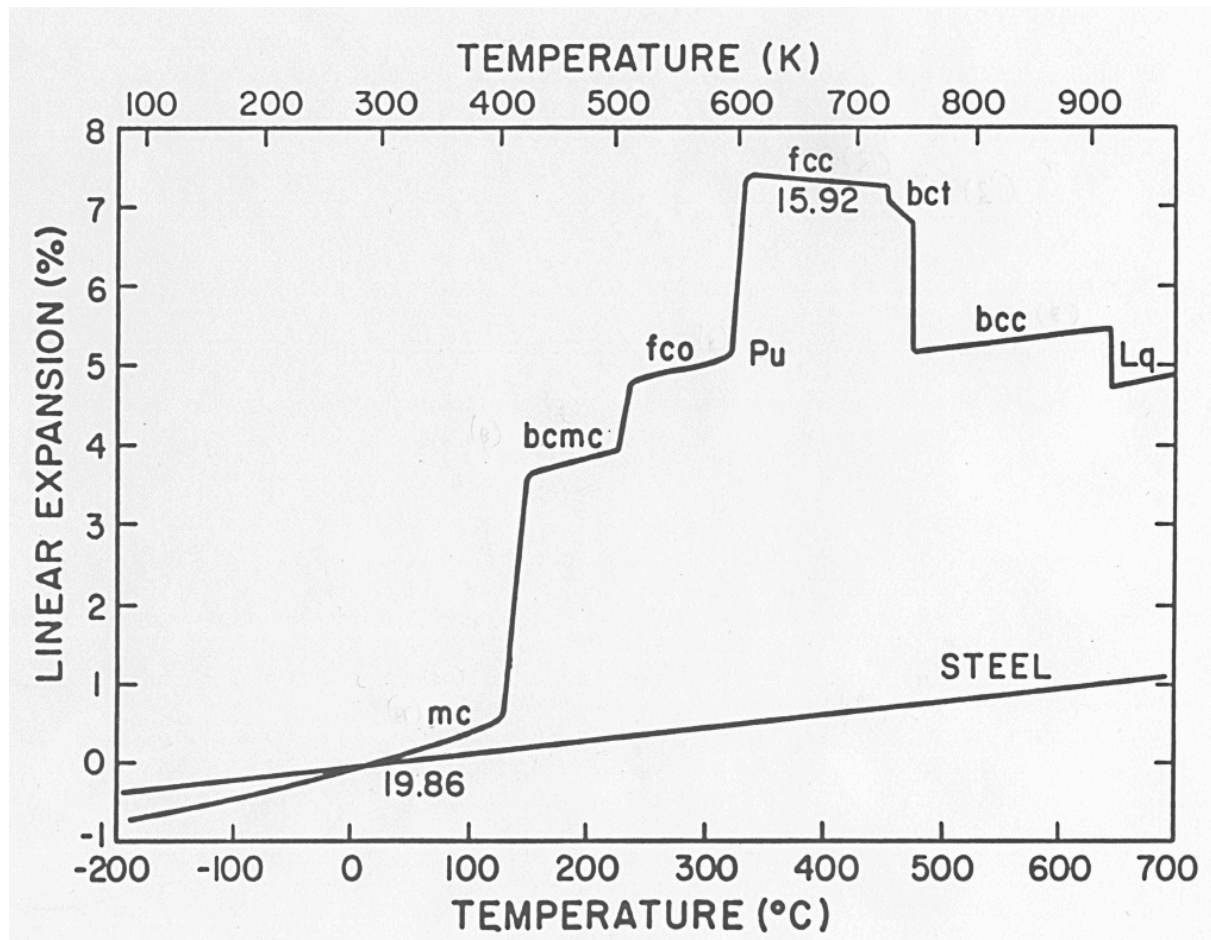
Past manufacturing process design...

- Was based on expert experience: trial-and-error, intuition
- “Works”, but is inefficient in time, energy, cost, and material

Can a simulation tool help?

- Reduce inefficiencies, provide insight, enable optimization?

Plutonium: Not Your Typical Industrial Metal





Project Customers, Collaborators

Los Alamos National Laboratory (LANL)

- Telluride project team members: applications, V&V, development
 - Group MST-6: gravity casting (Sigma, CMR foundries), welding
 - Group NMT-5: gravity casting (TA-55 foundry); MST-6: welding
- Group ESA-WMM: foam curing

Lawrence Livermore National Laboratory (LLNL)

- Die casting

Y-12 National Security Complex

- Gravity casting, vacuum arc re-melting (VAR)

Department of Defense (DOD) Missile Defense Agency (MDA)

- Droplet breakup

Universities

- CSM, UCI, Cornell, Memphis, UCLA, Iowa, University of Tennessee

Navy

- Naval Research Laboratory: welding modeling
- Naval Surface Warfare Center: droplet breakup



Casting Simulation: Impact List

Easiest: Provide qualitative insight

- About what has just happened; By answering basic “what if” questions

Easier: Guide macroscopic process parameters

- Gating & feeding rate & geometry, mold & charge preheat
- Mold geometry, cooling boundary conditions

Hard: Capturing microstructural effects

- Thermal history, alloy micro- & macro-segregation
- Correlation-based grain size and orientation predictions

Harder: Predictive, quantitative accuracy

- Part geometry: residual stress/distortion, porosity, cracks
- Material microstructure specifications: grain characterization

Hardest: Process optimization

- Based on part/material design criteria; Based on other criteria such as energy, cost, etc.

Casting Simulation Tools: Their Inabilities*

Inclusion formation (oxidation during pouring) and final location in casting

Hot tears

Mold penetration (with effects of coatings), sand defects

Leakers due to porosity , inclusions, etc.

Automated, optimized design of filling system

Simple geometric tools during initial concept development

Integration of casting models, nondestructive evaluation standards, and service performance analysis

Incorporation of variables such as mold hardness, metal chemistry variations

Prediction of final casting dimensions, taking into account shrinkages/stresses

Completing simulations in a realistic amount of wall clock time



Casting Simulation Tools: Their Abilities*

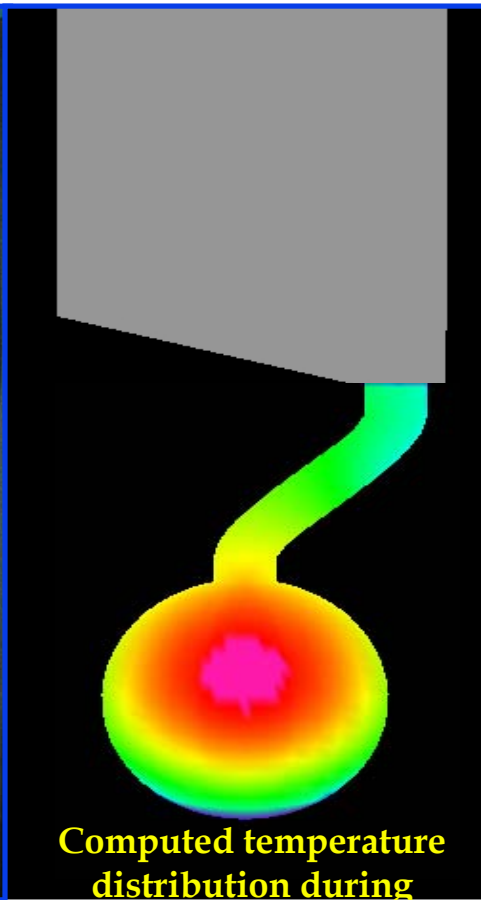
Heat transfer



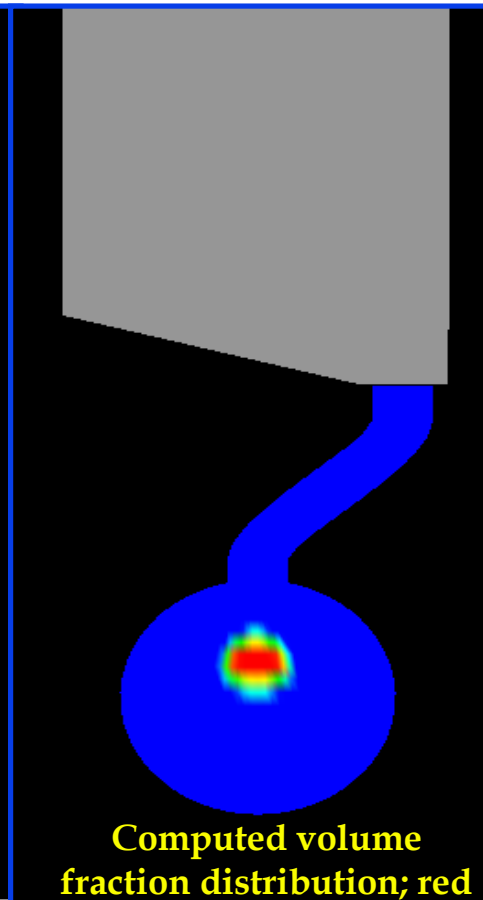
Application: Aluminum Ball Casting



Photo of aluminum ball casting showing porosity



Computed temperature distribution during solidification.



Computed volume fraction distribution; red is last to solidify

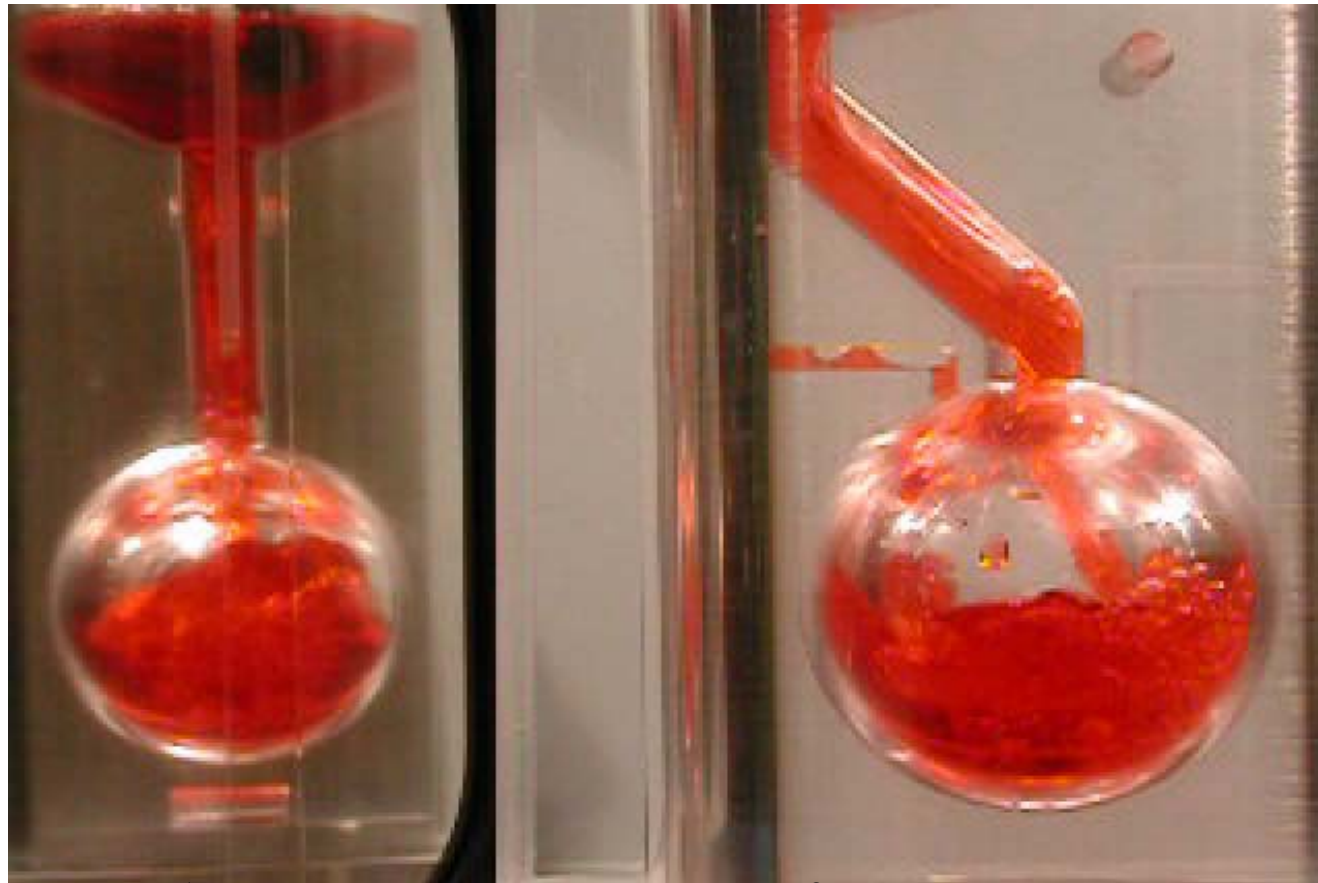
A ball is a difficult shape to cast since the last place to solidify will tend to be toward the center of the ball and will typically exhibit shrinkage porosity as shown.

Aluminum is a challenging metal to cast as it has 7% shrinkage on solidification.

Basic phase change heat transfer simulations really do make a difference! If ICs and BCs are known, the last point to solidify is an attainable simulation result

Experiments: Deniece Korzekwa
Calculations: Deniece Korzekwa

Mold Filling Experiments, Simulations



side view

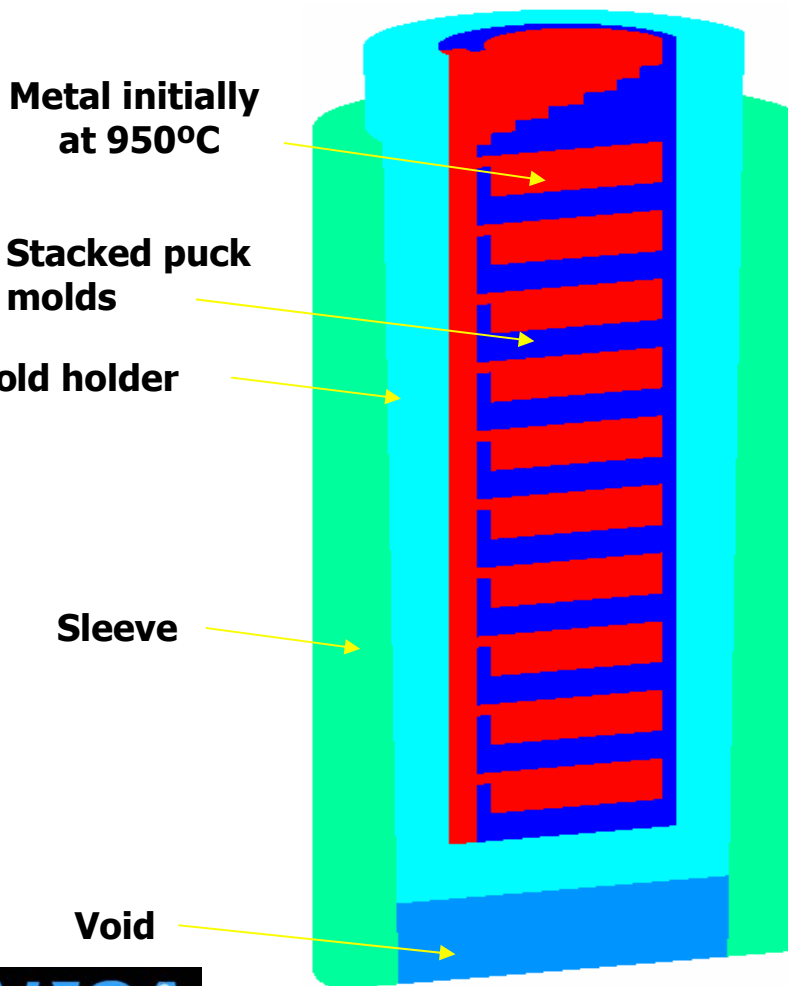
front view

water pouring into
lexan mold

Experiments:
Deniece Korzekwa

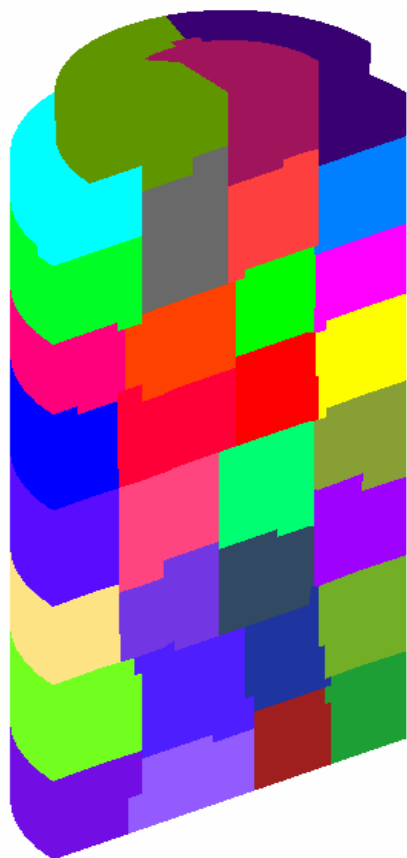
Calculations:
Markus Bussmann
Deniece Korzekwa
Kin Lam

Application: "Puck" Castings



- Mesh is half symmetry with 38364 cells and 41847 nodes.
- All mold parts are graphite
- Simulation starts with metal at 950°C and graphite at 650°C
- Simulations are used to predict the cooling rate and expected grain size.
- Typical CPU times (4-PE Compaq ES40)
 - 21264 Alpha processor (667 MHz)
 - 2 PEs - 318 μ s/cell/cycle
 - 4 PEs - 110 μ s/cell/cycle
- Experiments: Dave Olivas
- Calculations: Deniece Korzekwa

Application: “Puck” Castings



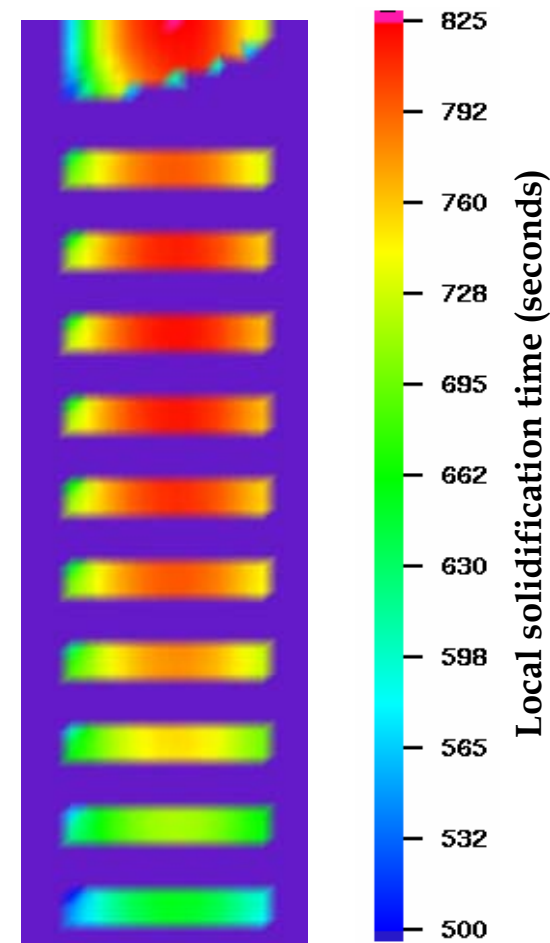
32 processor
decomposition

Local solidification time (time in mushy zone) and the cooling rate through the epsilon phase have been correlated to grain size and coring.

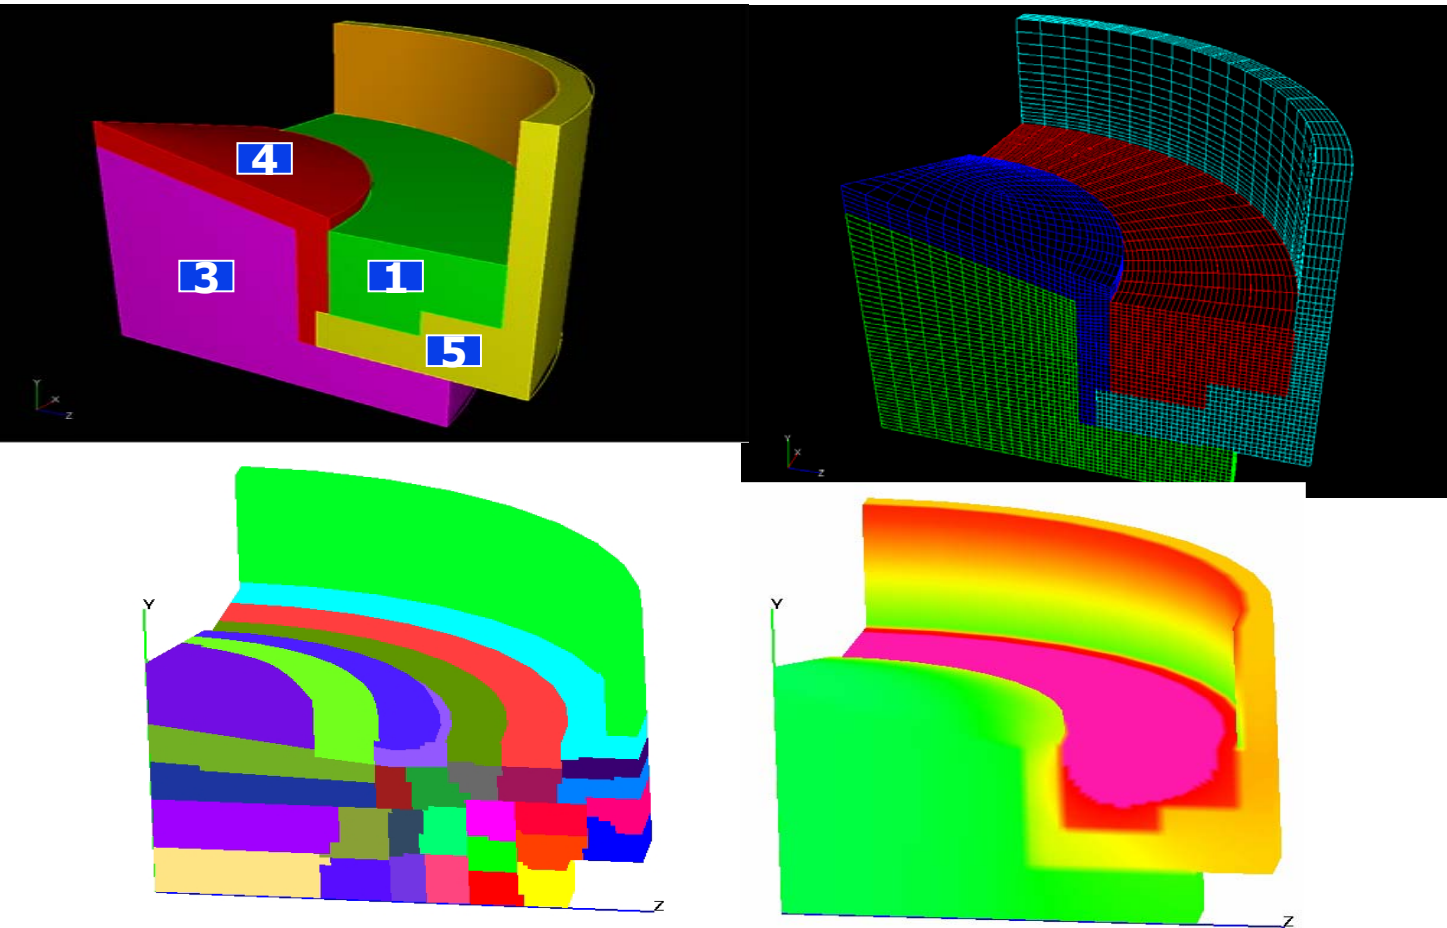
Simulations show variations between pucks as well as within each puck.

This information can be compared to other samples and be used in subsequent simulation codes.

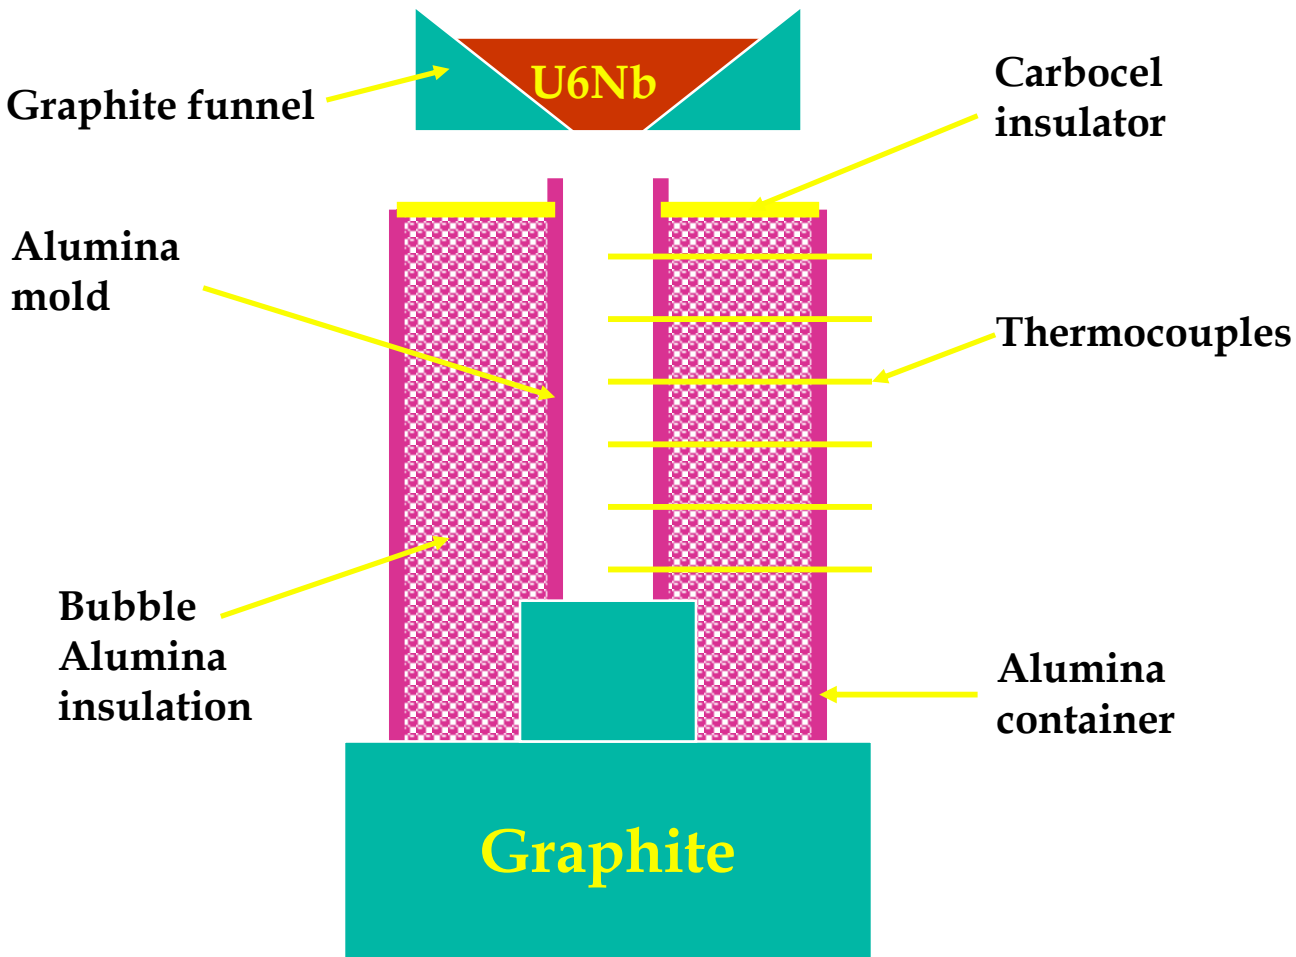
Calculations: Deniece Korzekwa



Application: SPR-III Reactor Ring Castings



Application: Directional Solidification Experiment



Unheated mold

U6Nb Temperature
1350-1475°C

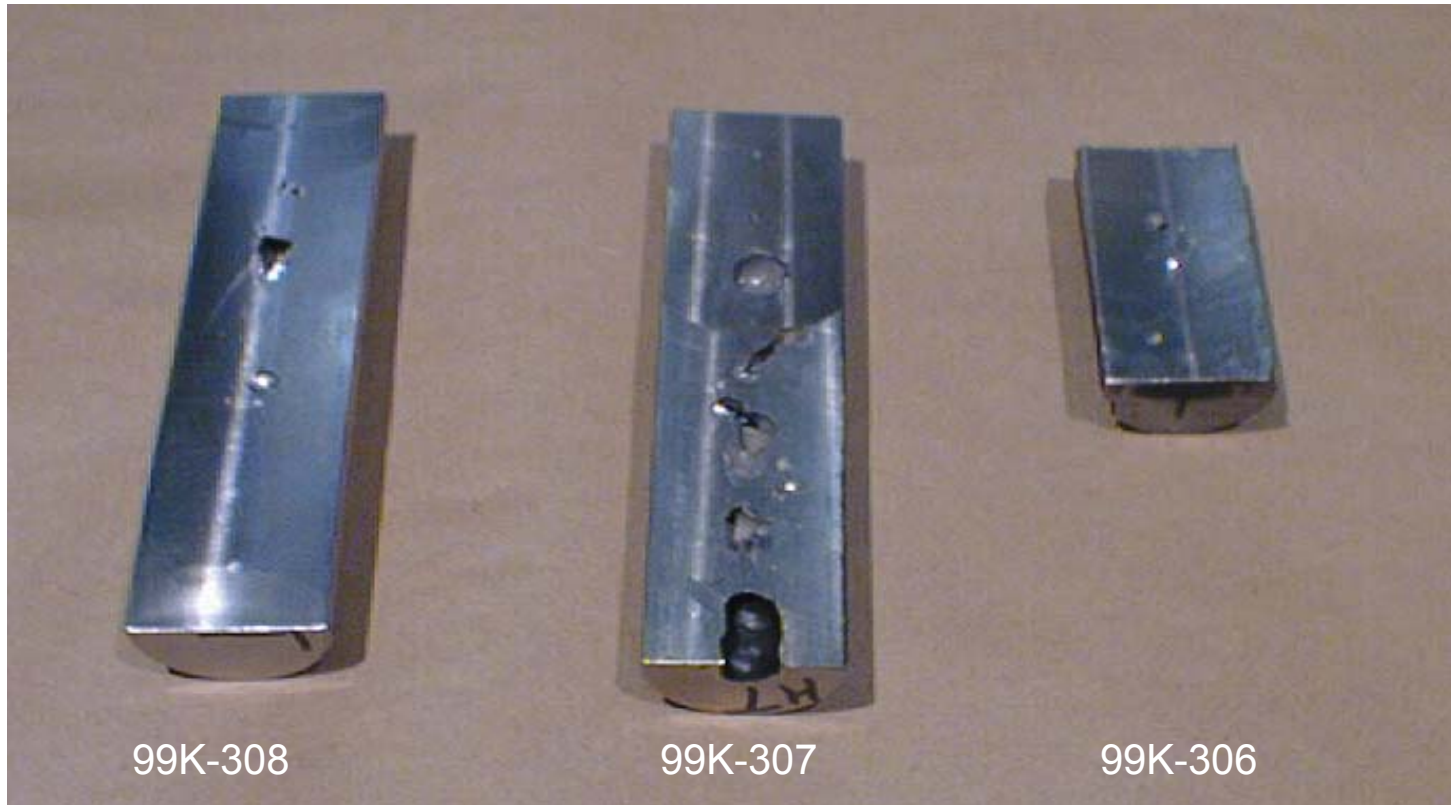
Type C and K
Thermocouples

Ideally unidirectional
solidification

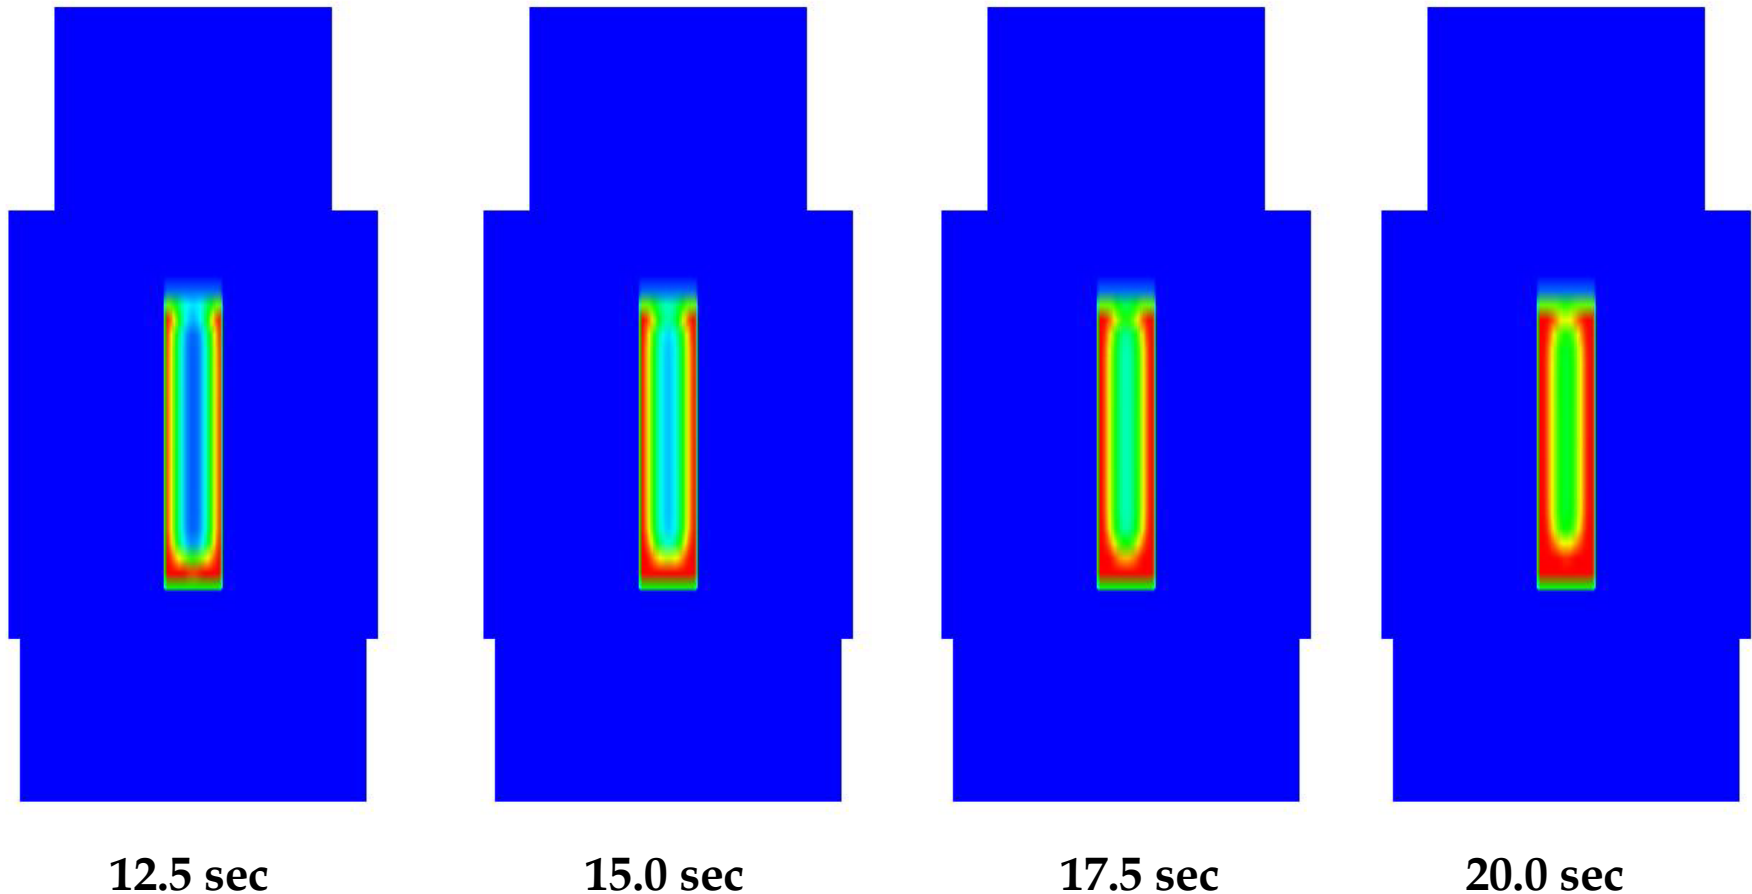
Experiments:
Deniece Korzekwa

Calculations:
Doug Kothe

Application: Directional Solidification Experiment

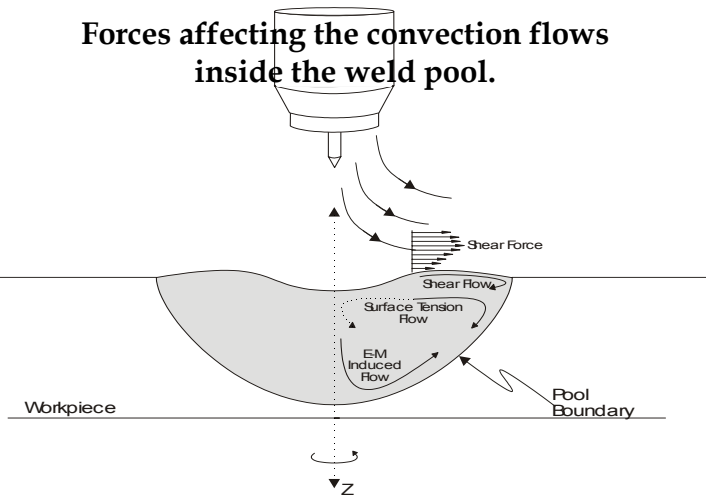


Simulation Results Non-Axial Solidification Front Propagation

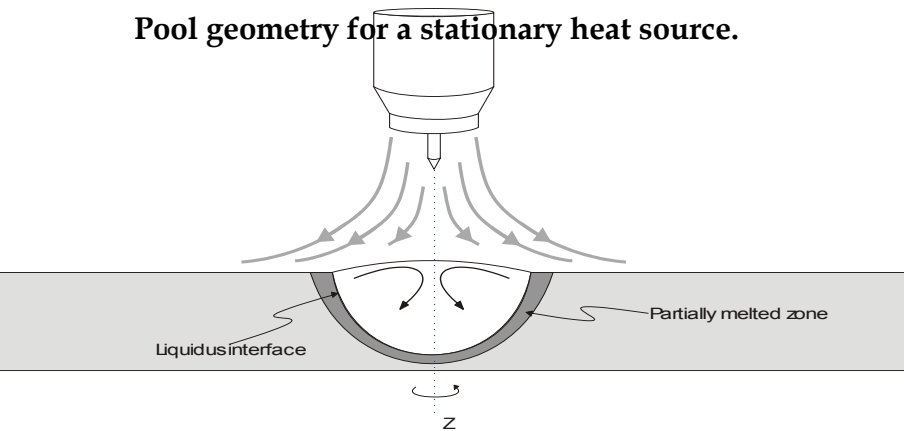


Welding Phenomena: Convection Flows and Traveling Heat Sources

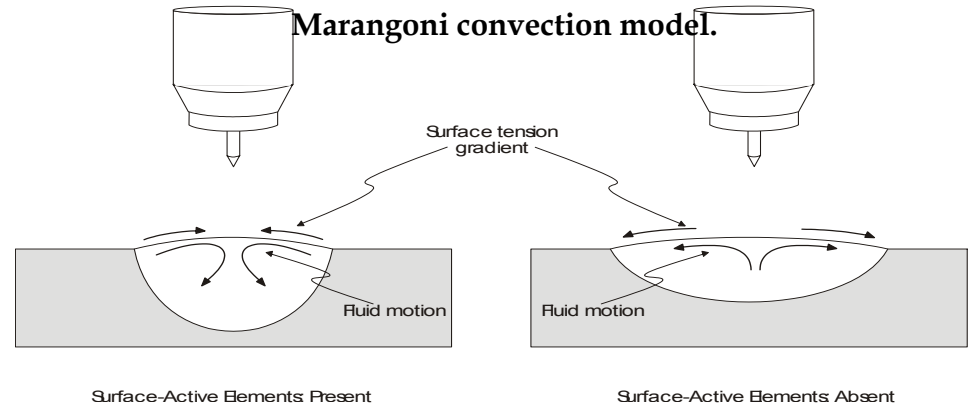
Forces affecting the convection flows inside the weld pool.



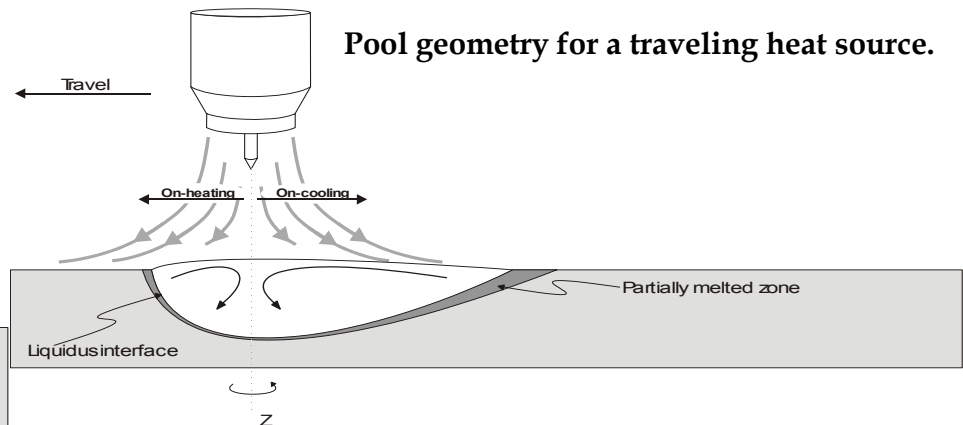
Pool geometry for a stationary heat source.



Marangoni convection model.

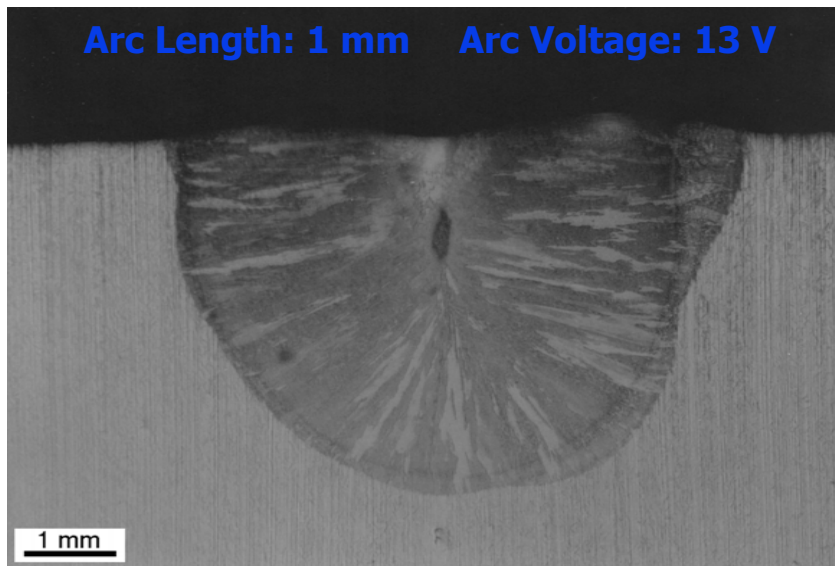


Pool geometry for a traveling heat source.



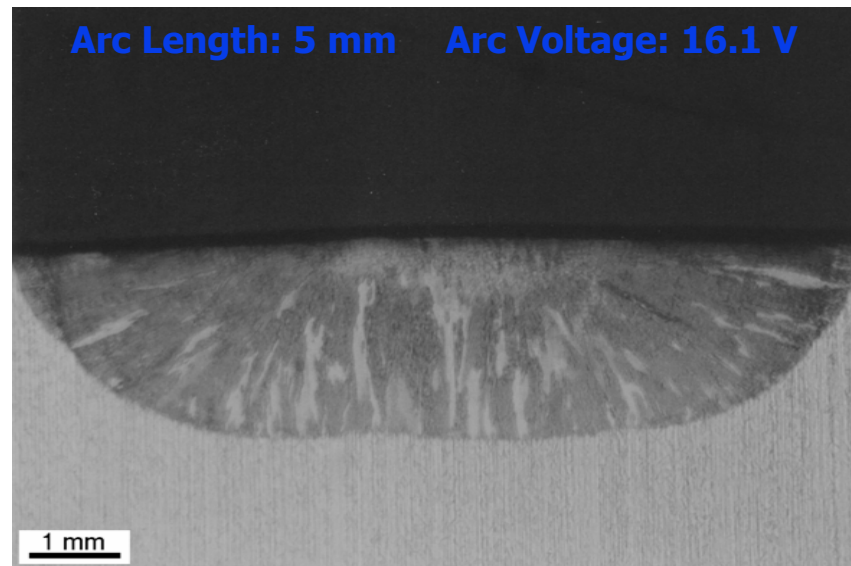
Useful Welding Models Must Predict Weld Pool Dimensions

LANL GTAW Spot Weld Experiments on SS304L



Width: 6.37 mm

Depth: 4.89 mm



Width: 9.02 mm

Depth: 2.60 mm

Constant parameters

♦ Arc current: 180 A

Beam on time: 5 s

♦ Shield gas: Argon

Flow rate: 25 scfh

Telluride Project Taxonomy

1995: The name “Telluride” was assigned to new software being prototyped at LANL for 3D free surface flow modeling

- Developed up to that time by D. Kothe, S.J. Mosso, and R. Ferrell
- The name was given by John Hopson (Kothe’s boss), a snow ski enthusiast who enjoys skiing at the Telluride Ski Resort in Telluride, CO
- No direct connection to the element “tellurium”
- Within a year, the Telluride name had a life of its own, applying not only to the software but the project as well
 - *A confusing point that persists today*

1999: The name “Truchas” was assigned to the subset of the Telluride software that is releasable to the general public

- Truchas is Spanish for trout, and also the name of a mountain peak in the Sangre de Christos mountain chain near Santa Fe, NM

Telluride Project Working Groups

Working Groups (WG) have improved project productivity since their establishment ~2 years ago

- Established to increase communication and to promote delegation of responsibility, accountability, ownership, and decision making
- Project lead is no longer the bottleneck

Each WG

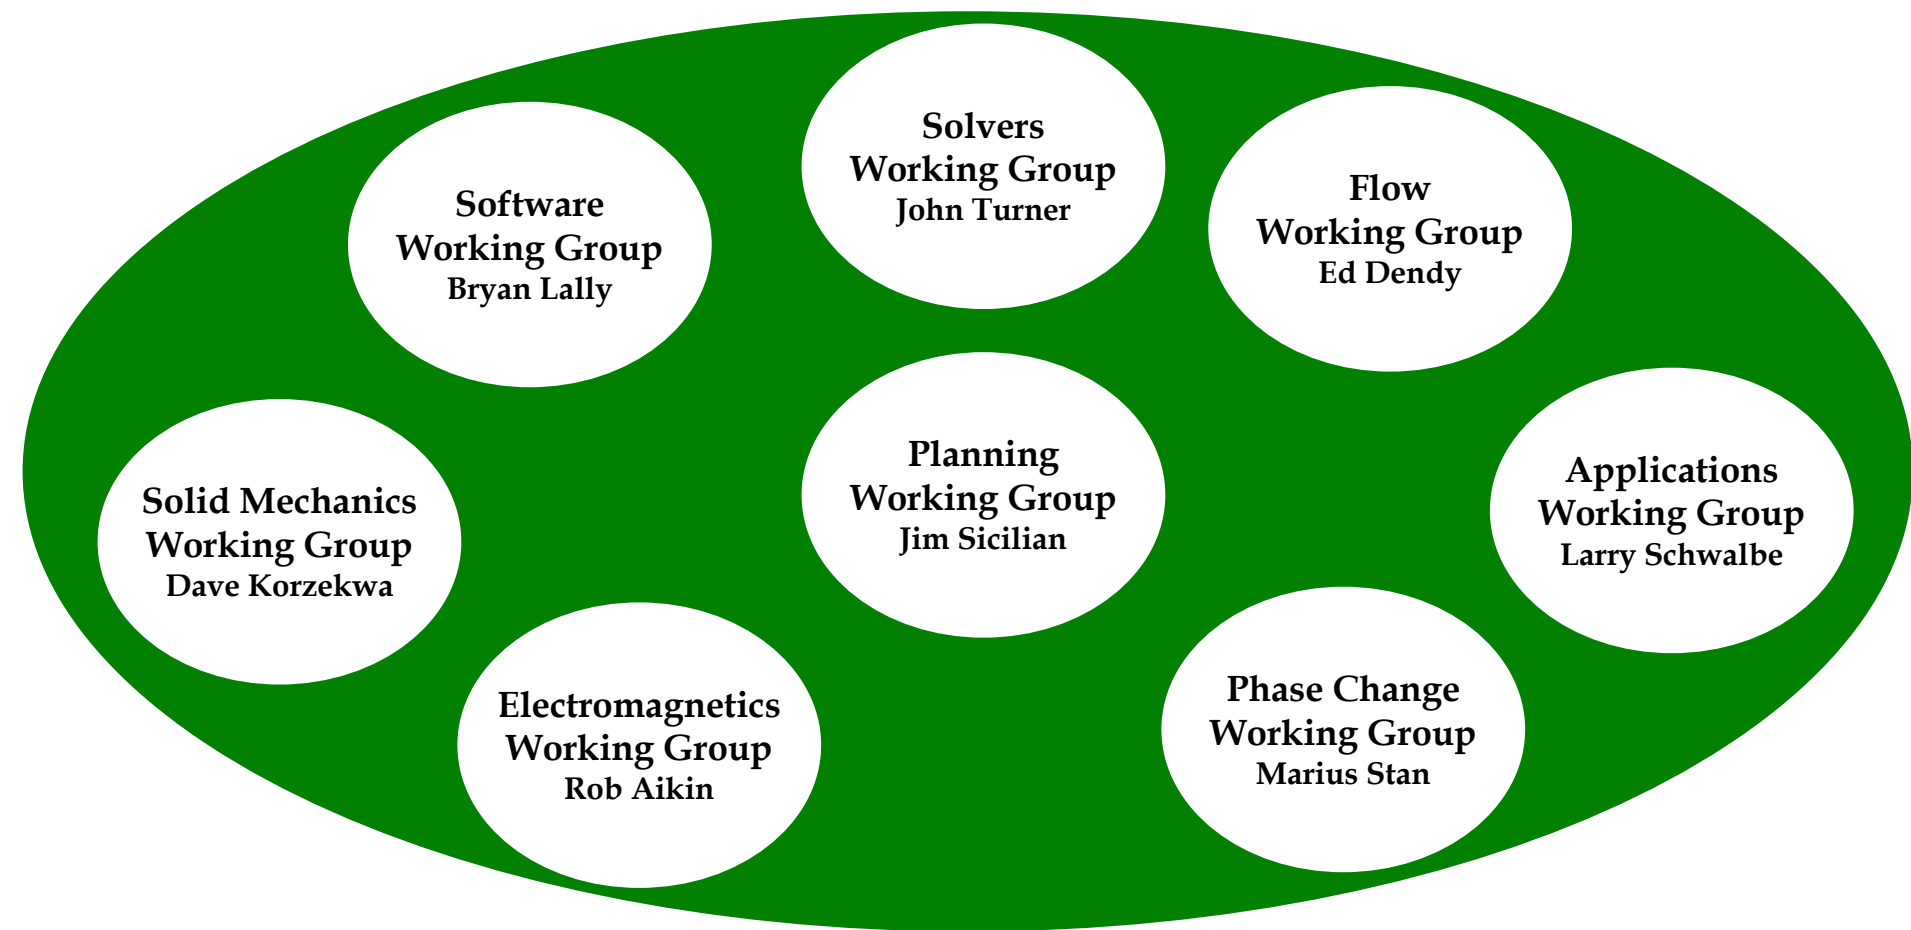
- is chaired by a lead who is a subject matter expert
- has 2-5 “card-carrying” member who meet bi-weekly for working and planning sessions
- is responsible for planning and executing its portion of the project

WGs foster leadership, ownership, pride

- Word has gotten around: Most of LANL Defense Programs now have WGs

WGs provide multiple conduits into and out of the project

The Telluride Project Is Subdivided Into Working Groups



The Telluride Project Team

Project Lead: Jim Sicilian

Team Members:

Rob Aikin

Sharen Cummins*

Robert Ferrell*

Deniece Korzekwa*

Bryan Lally

Jamaludin Mohd-Yusof

Sriram Swaminarayan

Jim Sicilian*

David Brown*

Ed Dendy

Marianne Francois

Doug Kothe*

Kin Lam

Larry Schwalbe

John Turner*

Rudy Henninger*

Markus Bussmann

Jeff Durachta

Dave Korzekwa*

Andrew Kuprat*

Sam Lambrakos

Marius Stan*

Matt Williams*

* presenting at this workshop

Many Others Have Contributed To The Telluride Project In The Past

LANL

- Brian Vanderheyden, Mark Schraad, Frank Harlow, Nely Padial, Dana Knoll, Vince Mousseau, Jerry Brock, Hilary Abhold, Jay Mosso, John Cerutti, Bill Rider, Richard LeSar, Mike Johnson, Laura Crotzer, Mark Taylor, Dan Hartman

Elsewhere

- Mike Steinzig, Damir Juric, Phyl Crandall, Christoph Beckermann, Anand Reddy, Larry DeChant, Jianzheng Guo, Sam Johnson, Jeff Marchetta, Austin Minnich, Michelle Murillo

Telluride Project Working Groups

Planning (Jim Sicilian): John Turner, Bryan Lally, Dave Korzekwa, Robert Ferrell

*Flow (Ed Dendy): Doug Kothe, Jim Sicilian, Jamal Moyd-Yusof, Sharen Cummins,
Matt Williams*

*Phase Change (Marius Stan): Sriram Swaminarayan, Andrew Kuprat, Doug Kothe,
Matt Williams, Kin Lam*

*Software (Bryan Lally): John Turner, Robert Ferrell, Sriram Swaminarayan,
Ed Dendy*

Solid Mechanics (Dave Korzekwa): Robert Ferrell, Doug Kothe

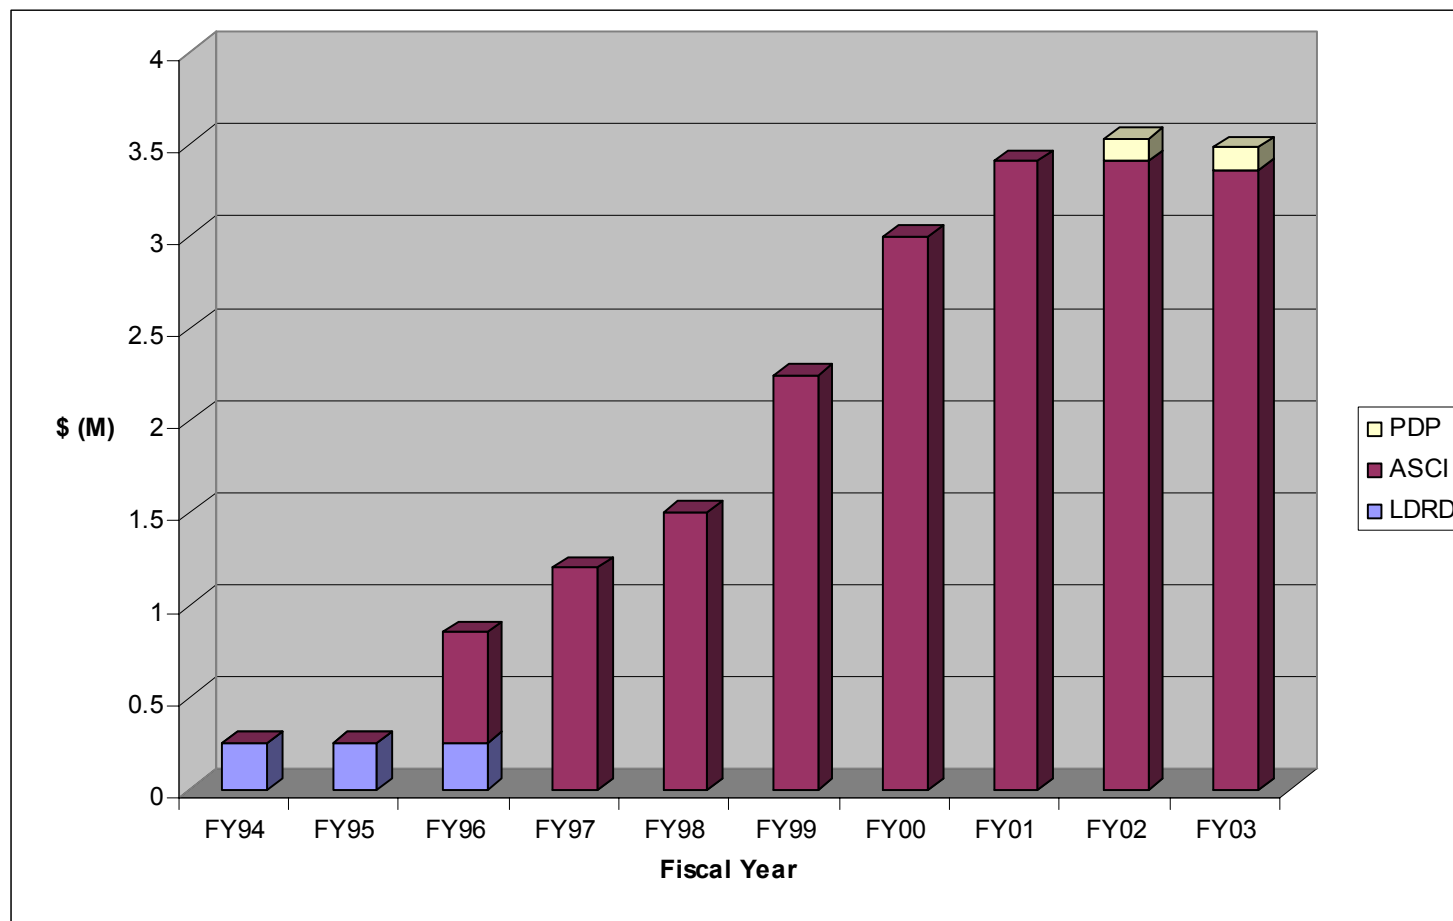
Electromagnetics (Rob Aikin): David Brown, Robert Ferrell

Solvers (John Turner): Bryan Lally, Robert Ferrell, Doug Kothe

Applications (Larry Schwalbe): Kin Lam, Deniece Korzekwa

V&V (Kin Lam)??

Telluride Project Budget History





The Role Of Our Project Software Relative To Commercial Software

Truchas is a LANL research code designed, implemented, tested, and applied specifically for simulating DOE manufacturing processes

- As mandated and supported by the DOE ASCI Program

Truchas is not a general purpose simulation tool

- That has been developed with commercialism in mind

Truchas is complementary to commercial software possessing related capabilities

- Truchas is inherently 3D, designed-in parallelism, high-fidelity (hence possibly higher risk) algorithms that are years (~5?) ahead of commodity use
- Truchas is “tuned” specifically for DOE manufacturing via appropriate models
- We have and will continue to use commercial software for simpler (2D) simulations

Through publications, workshops such as this, and software releases we will make available our models, algorithms, and software to commercial vendors

- One possible measure of success: a vendor cites using the “Truchas model/algorithm” for ...



The Telluride Project

Overview of the DOE ASCI Program

- which supports Telluride

Motivation, rationale, history

Physical models

Numerical algorithms

Software overview

Challenges and barriers

Physical Models: Overview of Requirements

Heat transfer

- Conductive, convective, inductive
- Radiative on boundaries

Phase change

- Liquid/solid, solid/solid transitions
- Liquid/vapor transitions?
- Pure materials, alloys having n species
- Arbitrarily-complex phase diagrams

Incompressible fluid flow

- Boussinesq approximation for thermal and solutal buoyancy
- Turbulence models for fill
- Porous media models in mush

Interface kinematics and dynamics

- Free surfaces, fluid/solid and fluid/fluid interfaces
- Surface tension and phase change forces

Thermo-mechanical material response

- Residual stress and distortion
- Elastic/plastic, creep, phase change

Quasi-static MHD

- Inductive (Joule) RF heating
- Magnetic (Lorentz) stirring

Material microstructure

- Nucleation and growth models
- Homogenization models

Physical Model for Flow: Incompressible Single-Fluid Navier Stokes

$$\frac{\partial \rho}{\partial t} + \mathbf{u} \cdot \nabla \rho = 0; \nabla \cdot \mathbf{u} = 0$$

More detail in Jim
Sicilian presentation

$$\frac{\partial}{\partial t} (\rho \mathbf{u}) + \nabla \cdot (\rho \mathbf{u} \mathbf{u}) - \nabla \cdot \boldsymbol{\tau} - \rho(C, T) \mathbf{g} - (\mathbf{F}_D + \mathbf{F}_s + \mathbf{F}_L) = 0$$

- All fluids move with a single center-of-mass velocity field (“single-field”)
- Newtonian stress tensor τ (future capability: visco-elastic stresses)
- Boussinesq approximation yields $\rho(C, T)$ dependence on gravity force
- Solid phases ($k=s$) are considered rigid: $\mathbf{u}_l = \mathbf{u}$; $\mathbf{u}_s = 0$
- Drag force \mathbf{F}_D for flow retardation in partial solid (mush) regions
- Surface tension (CSF) force \mathbf{F}_s models uneven interfacial attraction effects
- Lorentz force \mathbf{F}_L models stirring effects of magnetic fields (not available yet)

Physical Model for Heat Transfer: A Mixture Enthalpy Formulation

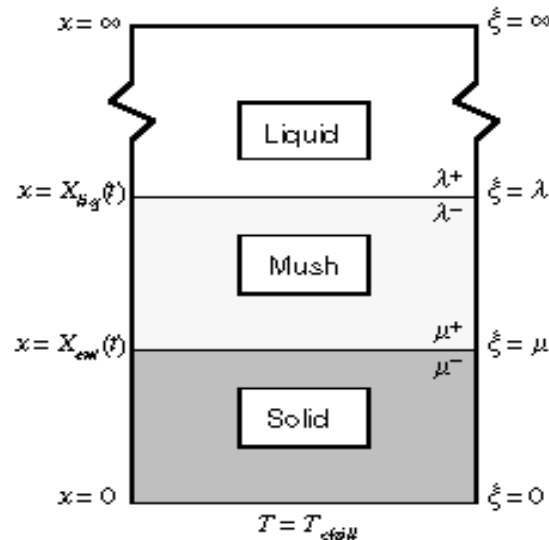
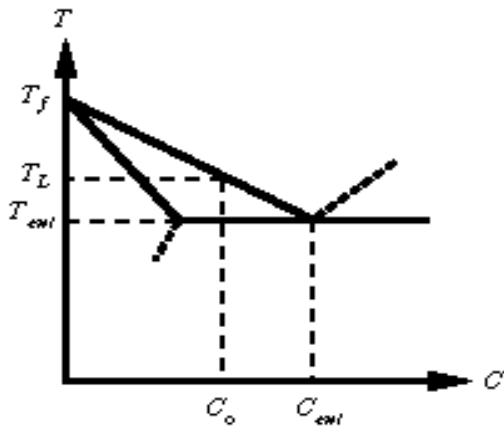
$$\frac{\partial}{\partial t} \langle \rho h \rangle + \nabla \cdot [\mathbf{u} \langle \rho h \rangle] - \nabla \cdot [k \nabla T(h)] - S(h) = 0$$

More detail in Marius
Stan presentation

$$\langle \rho h \rangle(T) = \sum_k \rho_k \varepsilon_k h_k(T); \quad h_k(T) = h_k^{\text{ref}} + \int_{T_k^{\text{ref}}}^T C_{p,k}(T) dT$$

- Thermodynamic equilibrium: $T_k = T$ for all phases k
- Enthalpy nonlinearities are present because of:
 - Nonlinear temperature dependence of $C_p(T)$ and $k(T)$
 - Jumps in reference latent heats h^{ref} at phase change boundaries
- Solid phases ($k=s$) are not allowed to flow: $u_l = u$; $u_s = 0$
- Phase fronts are *captured* (ε within a cell) rather than *tracked*
- A feature soon-to-be available: classical (JMAK) nucleation and growth model

Alloy Solidification: Micro-Segregation Model



Voller found similarity solutions for this system:

Int. J. Heat Mass Transfer, 40, 2869 (1997)

$$-\left[\rho h\right]+\nabla \cdot\left(\rho_l h_l \mathbf{u}_l\right)=\nabla \cdot\left(\nabla \kappa T\right) \quad \varepsilon_l \rho_l \frac{\partial C_l}{\partial t}+\varepsilon_l \rho_l \mathbf{u}_l \cdot \nabla C_l=\left(C_l-C_{si}\right) \frac{\partial}{\partial t}\left(\varepsilon_s \rho_s\right)+\frac{S_v \rho_s D_s}{l_s}\left(C_s-C_{si}\right)$$

$$\rho h]=\varepsilon_s \rho_s h_s+\left(1-\varepsilon_s\right) \rho_l h_l \\ =\int C_{ps} d T ; h_l=\int C_{pl} d T+L$$

$$\varepsilon_s \rho_s \frac{\partial C_s}{\partial t}=\left(C_{si}-C_s\right) \frac{\partial}{\partial t}\left(\varepsilon_s \rho_s\right)+\frac{S_v \rho_s D_s}{l_s}\left(C_{si}-C_s\right)$$

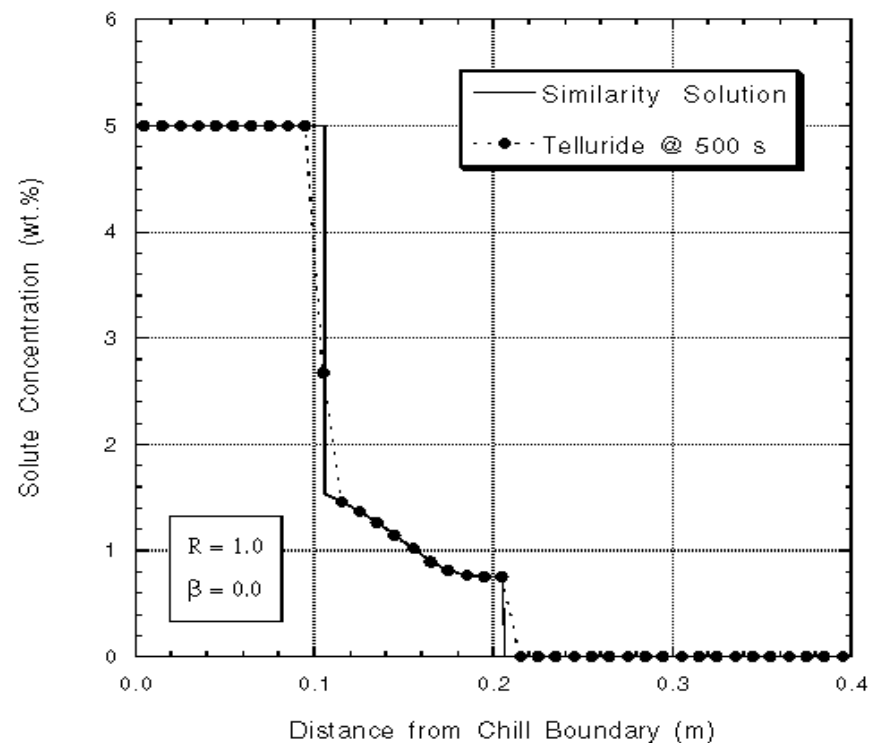
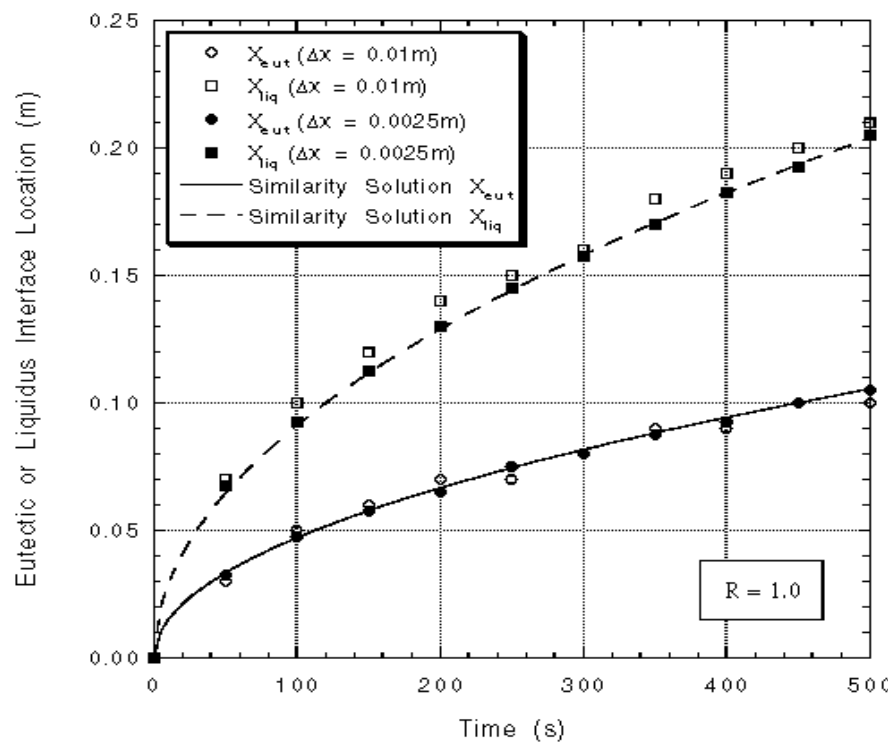
$$T=T_m+m_l C_{li} ; C_{si}=k C_{li}$$

alloy species transport and phase diagram

mixture enthalpy transport



Alloy Solidification Verification



Verification studies performed by Kin Lam

Physical Model for Solid Material Response

Quasi-static Thermal Elasticity

$$\frac{\partial \sigma_{ij}}{\partial x_j} = 0; \quad \sigma_{ij} = \lambda \varepsilon_{kk} \delta_{ij} + 2\mu \varepsilon_{ij} - (3\lambda + 2\mu) \alpha (T - T_0) \delta_{ij}$$

$$\varepsilon_{ij} = \frac{1}{2} \left(\frac{\partial d_i}{\partial x_j} + \frac{\partial d_j}{\partial x_i} \right)$$

More detail in Dave
Korzekwa presentation

$$\frac{\partial}{\partial x_i} \left(\lambda \frac{\partial d_k}{\partial x_k} \right) + \frac{\partial}{\partial x_j} \left[\mu \left(\frac{\partial d_i}{\partial x_j} + \frac{\partial d_j}{\partial x_i} \right) \right] - \frac{\partial}{\partial x_i} \left[(3\lambda + 2\mu) \alpha (T - T_0) \right] = 0$$

- Displacements d_i are assumed to be “small” (Eulerian frame)
- Material constants α , λ , and μ are in general temperature dependent
- Can be easily expanded to include plasticity, viscoplasticity

Physical Model for Electromagnetics

Defer this to David Brown, Robert Ferrell, and Rob Aikin

- And their collaborators Misha Shashkov (LANL Group T-7) and Neil Carlson (Motorola)

The Telluride Project

Overview of the DOE ASCI Program

- which supports Telluride

Motivation, rationale, history

Physical models

Numerical algorithms

Software overview

Challenges and barriers

Numerical Algorithms: Overview

Unstructured mesh domain partitioning (fixed, Eulerian reference frame for now)

- Each cell is a logical cube mapped into degenerate hexahedra: hexes, tets, prisms, pyramids

Cell-centered collocated finite volume discretizations

- Nearest neighbor connectivity; formally 2nd order in space
- Embrace weighted least squares for function reconstruction and some matrix operators (Laplacian)

Phase change heat transfer

- Implicit, nonlinearly-consistent NK-based enthalpy method
- Locally nonlinear (point) systems for alloy solidification

Flow: semi-implicit, fractional step projection method

- Moving to fully implicit NK-based scheme as an option
- Design constraint: complex topology interfacial flows with density ratios of 1 to infinity (void) across interface

Solid Mechanics: node-based CV-FEM approach

Electromagnetics: edge, face-based staggering

Multi-level-preconditioned Krylov-based linear solution methods

The Flow Algorithm: Highlights

Finite volume, cell-centered projection method

- All permanent variables are collocated at cell centroids
- “MAC” projection of face-centered velocities to enforce discrete solenoidality
- Rhie-Chow-like pressure gradient treatment for face velocities because of a pressure correction projection

Optional implicit Newtonian viscosity treatment

New 2nd-order “corner-coupled” flux-limited advection algorithm for smooth and non-smooth data

Void model is available, with a new void closure model under development

Novel, consistent treatment of variable (momentum, enthalpy, species) advection within interface cells (one result: enables arbitrarily high density ratios)

Piecewise-linear volume tracking algorithm for interface kinematics

Continuum surface force (CSF) basis for interface dynamics (surface tension)

Phase Change Heat Transfer Algorithm: Highlights

Global Newton-Krylov (NK) nonlinear solution algorithm

- With optional damping and two-level preconditioning
- Allows arbitrarily high time step, yet with a novel adaptive time step control

Local Newton-Raphson nonlinear solution for cells undergoing phase change

- Written in a completely general way to handle arbitrarily complex phase diagrams and $T(H)$ functions

Variety of BC types supported

Nonlinear property (C_p , k , etc.) relations easily supported

Potentiality for increased fidelity near interfaces

Chemical reaction models (exo- or endothermic) easily supported

Support for surface and volumetric heat source/sinks

Solid Mechanics Algorithm: Highlights

Node-centered CV-FEM discretization

- With two different choices for order of spatial accuracy
- Follows, to some extent, the Physica implementation
- Created algorithm basis useful for other node-centered schemes

Novel preconditioning matrix and preconditioning solution method

Potentiality for allowing mesh to move and distort in a Lagrangian sense

Electromagnetics Algorithm: Highlights

Finite element scheme

Clever use of edge- and face-based staggering in order to rigorously preserve important vector identities

Use of “Whitney” element basis functions

Preservation of charge and energy via a clever linear solution derived from material constitutive equations

Time-implicit discretization in order to integrate over fast-moving (light speed) waves

Linear and Nonlinear Solution Methods: Highlights

New UbikSolve Krylov-based library

- With matrix-free matvec support, preconditioning call-back support
- Embraces FGMRES to allow more flexible preconditioning

MG-like two-level preconditioning even when running in serial

- Usually gives substantial speedups relative to more traditional preconditioning

A beginning of a “black-box Newton-Krylov” capability for arbitrary nonlinear functions $F(x)$

Parallel Algorithms and Software

Our algorithms and software were designed with efficient, portable parallelism as a requirement

- We have relied on the expertise and guidance of parallel software expert Robert Ferrell
- One outcome: inception, design, and use of Robert's PSGLib library for "hidden parallelism"

Embrace the single-program, multiple-data (SPMD) paradigm

- Using MPI as the low-level communication library

Notable attributes

- Parallelism is hidden from the user
 - *Just use a different executable and tweak a few optional run-time input variables*
 - *Our output files are the same irregardless of the hardware platform used!*
- Parallelism is virtually hidden from the developer
 - *Except for a few "rules of engagement" with regard to communication primitives, etc.*

The Telluride Project

Overview of the DOE ASCI Program

- which supports Telluride

Motivation, rationale, history

Physical models

Numerical algorithms

Software overview

Challenges and barriers

Telluride Project Software Overview

Truchas is the principal software product

- Approximately 100K lines (including packages) of source code primarily written in Fortran 90/95

Yet there are other state-of-the art software packages used and required by Truchas

- PGSLib: Parallel Gather-Scatter Library (Robert Ferrell)
 - *Performs all necessary MPI-based parallel functions and communication*
 - *MPI function calls written in C; Fortran 95 interface provided*
- UbikSolve: Krylov-based linear solution library (John Turner)
 - *Returns solutions to $Ax=b$; comes with a host of preconditioning options*
 - *Written in clean, portable Fortran 90/95*
- Chaco: domain decomposition library available from SNL (Bruce Hendrickson)
 - *Soon to be replaced by SNL's new Zoltan package*
- All are also freely available as standalone libraries (separate from Truchas)

Telluride Project Software Practices and Processes

Many project team members have been formally trained in modern software engineering practices and software project management

- We aspire to embrace practices such as configuration management, proven development cycles (requirements, design, implementation, testing, support), rigorous V&V, unit testing, component-based software architecture and design, exhaustive documentation, software design and source code reviews, etc.
- But it has been and will continue to be a long, slow uphill battle

Our project software has already been audited by the DOE in 1999

- We are due for additional audits:
 - An internal (LANL CCS Division) audit within the next year
 - Another external (DOE) audit in 2004

Bottom line: we must take the software aspect of this project seriously!

The Telluride Project

Overview of the DOE ASCI Program

- which supports Telluride

Motivation, rationale, history

Physical models

Numerical algorithms

Software overview

Challenges and barriers

Challenges and Barriers: Physics

Coupling the micro and macro length scales

- Common theme in other fields (turbulence, climate modeling)
- Averaged models inspired by DNS data sets and homogenization theories (*mushy zone transport*)
- Can micro and macro simulations be coupled simultaneously?

Phase change heat transfer

- Non-equilibrium kinetics; microstructural nucleation and growth models
- Solid state transitions and their role in stress/distortion buildup
- Generalization to arbitrary equilibrium phase diagram relationships
- Boundary condition models for internal gaps

Material response

- Creep and hot tearing models
- Microporosity growth models
- Solid state transition effects (TRIP models)

Challenges and Barriers: Algorithms

Finite-width interface representation on fixed meshes

- Tracking fluid interfaces in the presence of rigid solids (Ex: molds, solidified material)
- Subgrid models ("triple points"), transitioning to and from a homogeneous mix model

Discrete operations (Div, Curl, and all that) around interfaces where physics change

- Imposition of constraints, reduction of order

Linear/Nonlinear solution algorithms

- Matrix-free Krylov and Newton-Krylov methods are great
 - A lot still depends on physics-based preconditioning; How much can be pushed into "black boxes"?
- When to couple and when not to couple? Cost/benefit?

Boundary condition treatments

Balancing robustness with accuracy; cost/benefit?

Coupling all the physics! How to do it and how tightly coupled?

Challenges and Barriers: Computer Science and Software Engineering

Single PE execution time; >80% of cycles moving data

- We would like to be 5-10X faster than currently at present!*

Efficient software processes and practices

- Frameworks, components, environments, in-situ prototyping*

Finding the right mix of efforts devoted to software engineering and physics models & algorithms

Object-oriented languages: evolution or revolution?

Balancing portability with performance

Rigorous and automated unit testing and V&V

Project Highlights

I have my own opinion, but hopefully this will become apparent by the end of the workshop

Contact Us

Email

- Software “support”: *telluride-support@lanl.gov*
- General information: *telluride-info@lanl.gov*
- Individual team members

Phone

- Project Lead Jim Sicilian (505-665-6827 or 505-667-7029)

Web

- Project web site is *www.lanl.gov/telluride*

Workshop “Ground Rules”

Feel free to interrupt during our presentations with your questions and comments

- We think an informal atmosphere will be more productive

Grab us for private discussions during breaks

Jot down your comments, concerns, and questions

- So it will be easy for us to receive your written feedback

Brainstorm how you might work with us

- In a way that helps us and helps you